arXiv:1405.7081v3 [q-bio.BM] 16 Dec 2014

Protein sectors: statistical coupling analysis versus conservation

Tiberiu Teşileanu*[1,2], Lucy J. Colwell[1,3], and Stanislas Leibler[1,4]

[1]The Simons Center for Systems Biology and The School of Natural Sciences, Institute for Advanced Study, Einstein Drive, Princeton, NJ 08540, USA
[2]Initiative for the Theoretical Sciences, CUNY Graduate Center, 365 Fifth Avenue, New York, NY 10016, USA
[3]Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB3 0WA, UK
[4]Center for Studies in Physics and Biology and Laboratory of Living Matter, The Rockefeller University, 1230 York Avenue, New York, NY 10065, USA

Abstract

Statistical coupling analysis (SCA) is a method for analyzing multiple sequence alignments that was used to identify groups of coevolving residues termed "sectors". The method applies spectral analysis to a matrix obtained by combining correlation information with sequence conservation. It has been asserted that the protein sectors identified by SCA are functionally significant, with different sectors controlling different biochemical properties of the protein. Here we reconsider the available experimental data and note that it involves almost exclusively proteins with a single sector. We show that in this case sequence conservation is the dominating factor in SCA, and can alone be used to make statistically equivalent functional predictions. Therefore, we suggest shifting the experimental focus to proteins for which SCA identifies several sectors. Correlations in protein alignments, which have been shown to be informative in a number of independent studies, would then be less dominated by sequence conservation.

*e-mail: ttesileanu@gmail.com

Contents

1 Introduction **3**

2 Results **5**
2.1 PDZ domains 7
2.2 Dihydrofolate reductase 10
2.3 Voltage-sensing domains of K^+ channels 12
2.4 *E. coli lac* repressor 12
2.5 Top eigenmode of the SCA matrix 12

3 Discussion **16**
3.1 Proteins with multiple SCA sectors 16
3.2 Conclusions 17

4 Methods **18**
4.1 Sequence alignments 18
4.2 Statistical coupling analysis 18
4.3 Sequence conservation 19

5 Acknowledgments **19**

A Calculating sequence covariance **20**

B Statistical coupling analysis **22**

C Alignments **24**

D Details about the DHFR analysis **24**

E Alanine scans **26**

F Diagonal of SCA matrix instead of conservation **26**

G Multiple sectors **28**

H Top eigenmode of SCA matrix—some details **30**

I Data availability **31**

J Differences from previous version **32**

1 Introduction

A fundamental question in biology is the relation between the amino acid sequence of a protein and its function and three-dimensional structure. Given the rapid growth in the sequence data available from many organisms, it has become possible to use statistical sequence analysis to approach this question. Based on sequence similarity, protein sequences can be grouped into families thought to share common ancestry; the proteins in such a family typically perform related functions and fold into similar structures [1, 2]. It has been shown in many studies that a statistical analysis of a multiple sequence alignment (MSA) corresponding to a given protein family can be used to find amino acids that control different aspects of a protein's function or structure.

A basic statistical quantity that can be calculated for a multiple sequence alignment is the distribution of amino acids at each site. In particular, the level of sequence conservation at each site is of biological relevance, since it is expected that conservation is low in the absence of selective pressures. For this reason, conservation has long been used to predict which parts of a protein are most likely to be functionally significant [3–7].

More recently, the availability of large sets of protein sequences has made it possible to also estimate higher-order statistics, such as the correlations between the amino acids found at each pair of sequence positions. In a number of examples, these statistics have been shown to contain information about the structure and function of proteins [8–12]. One way in which pairwise correlations might arise is for a deleterious mutation at a given position to be compensated by a mutation at a different position. This can yield a scenario in which the two individual mutations are relatively rare, but the combination of both is common in natural proteins.

Statistical coupling analysis (SCA) was introduced by Lockless and Ranganathan in 1999 as a way to infer energetic interactions within a protein from a statistical analysis of a multiple sequence alignment [13]. The authors compared the statistics of an alignment of PDZ domain sequences to measurements of the binding affinity between a particular member of the alignment ($\mathrm{PSD95^{pdz3}}$) and its cognate ligand. The statistical analysis assumed that the frequencies of mutations obey a Boltzmann distribution as a function of binding free energy, allowing estimation of the binding affinity by $\Delta G_i \sim \log f_i$, where f_i is the frequency of an amino acid type at a given site in the alignment. By conditioning on amino acid type at a second site, they calculated the amount by which the effect of a mutation at one site changed depending on the amino acid present at the second site: $\Delta\Delta G_{i|j} = \Delta G_{i|j} - \Delta G_i \equiv \Delta\Delta G^{\mathrm{stat}}$. This gave an estimate for the effective coupling between the sites.

The assumptions behind the original formulation of SCA are likely to be violated, since the selective pressures acting on a protein are more complex than simply maximizing binding to a ligand. Despite this, the method seemed to be effective. In the original paper [13], mutant cycle analysis was used to measure $\Delta\Delta G^{\mathrm{binding}}$, the amount by which the effect of a given mutation on ligand binding affinity of $\mathrm{PSD95^{pdz3}}$ changes when the mutation occurs on a background containing a second mutation. This can be written as $\Delta\Delta G^{\mathrm{binding}} = \Delta G_{i|j} - \Delta G_i$, where now ΔG represents a change in the physical free energy as opposed to a statistical construct. The quantity $\Delta\Delta G^{\mathrm{binding}}$ was observed to be well correlated with the statistically-calculated $\Delta\Delta G^{\mathrm{stat}}$. The set of residues identified by SCA to be coupled with a particular site known to be important for binding specificity of the PDZ domain was found to physically connect distal functional sites of the protein [13]. This led to the suggestion that these residues may mediate an allosteric response. Experimental evidence later showed that indeed some of the residues identified by SCA participate in allostery [14–17]. Moreover, in a different study, a large fraction of the artificial WW domains built by conserving the pattern of statistical couplings calculated by SCA were observed to be functional, while sequences built to conserve single-site statistics alone were not [8, 9].

Table 1. The different meanings that can be associated with protein sectors

Interpretation of sector	Signature	Possible ways of quantitative exploration
statistical	clusters of correlated mutations in MSA	statistical analysis of MSA
evolutionary	maintains identity under evolutionary dynamics	artificial evolution experiments
structural	distinct physical properties compared to surroundings	NMR, room-temperature X-ray crystallography, MD simulations
functional	altering sector positions changes functional properties	mutagenesis studies

Motivated by these observations, Halabi et al. reformulated SCA in purely statistical terms, avoiding the assumptions related to energetics [18]. The reformulation amounted to a particular way of combining correlations with conservation. The basic idea was to multiply each element of the covariance matrix C_{ij} by a product $\phi_i\phi_j$, yielding the "SCA matrix" $\tilde{C}_{ij} = \phi_i\phi_j C_{ij}$. The "positional weights" ϕ_i were a function of the frequency f_i of the most prevalent amino acid at each position, and were roughly given by $\phi_i \sim \log[f_i/(1-f_i)]$. This particular form was chosen to reproduce the results from the original formulation of SCA [18, 19]. In subsequent work regarding SCA, several variations on this basic idea were used; all of these yield similar though not identical results and are described more precisely in sections 4.2 and B.

Running either the original or the reformulated analysis on several examples [8, 9, 13, 16–18, 20], it was noticed that the resulting SCA matrix had an approximate block structure. In analogy to previous work in finance, Halabi et al. analyzed this structure by looking at the top eigenvectors of the SCA matrix [18, 21]. The corresponding groups of residues were called "protein sectors" because similar clusters observed in the correlations of stock prices were found to correspond to financial sectors. Experiments found that mutating residues in distinct sectors specifically affected different phenotypes of the protein [18], leading to the suggestion that each SCA sector might comprise a group of amino acids that control a particular phenotype.

It is important to note that there are several subtly different meanings that have been attributed to protein sectors (see Table 1). The description outlined above defines protein sectors as the results of a statistical analysis of a multiple sequence alignment. This definition depends on the statistical method employed; it would, for example, depend on the choice of positional weights in the case of SCA, or on the precise thresholds and methods used for clustering. To distinguish this from other meanings, we will call these *statistical sectors* (or *SCA sectors* when the statistical method is SCA).

The sectors identified by SCA have also been given an *evolutionary* interpretation [18, 20, 22], based on the fact that they are defined as groups of residues whose mutations are correlated in an alignment, and the sequences in the alignment are likely to be evolutionarily related. However, this argument is insufficient to prove the evolutionary nature of the statistical sectors, given that their precise composition is dependent on the statistical method employed [23]. Thus it is difficult to assess to what extent the sector's composition is actually related to the evolutionary process itself, as opposed to the choice of the statistical method. Strikingly, Halabi et al. showed that for an alignment of serine proteases, one of the sectors can be used to distinguish between vertebrates and invertebrates, suggesting that indeed an evolutionary interpretation may be appropriate [18]. However, before concluding that in general SCA sectors have an evolutionary interpretation, it would be important to extend these studies to different alignments. An alternative, more direct,

approach would be to perform artificial evolution experiments to check whether the SCA sectors maintain their integrity under strong selection, or whether new sectors can be created in this way. In addition, such experiments would provide data on the evolutionary dynamics of proteins, and thus help to define more precisely the notion of evolutionary sectors.

Another surprising property of the groups of residues identified by SCA is that they usually form contiguous structures in the folded protein, although they are not contiguous in sequence [18, 20, 22, 24, 25]. This suggests the notion of *structural sectors*, groups of residues having different physical properties compared to their surroundings. An experimental test for such inhomogeneities inside proteins could employ NMR spectroscopy to follow the dynamics of specific atoms while the protein is undergoing conformational change [14, 26]. In addition, analyzing room-temperature X-ray diffraction data could shed light on residues with coupled mobility or increased fluctuations in an ensemble of structures [27, 28] (Doeke Hekstra, personal communication). Alternatively, this kind of experiments could be done *in silico* using for example molecular dynamics simulations to identify correlated motions in the protein [29] (Olivier Rivoire, personal communication).

Finally, as mentioned above, a number of mutational studies have suggested yet another interpretation of the sectors identified by SCA as *functional sectors*, groups of amino acids that cooperate to control certain phenotypic traits of a protein, such as binding affinity [13, 18, 20, 25], denaturation temperature [18], or allosteric behavior [14–17, 24]. It is this aspect of the sectors that has been most emphasized in the literature.

In the language we just introduced, we can say that there is some data suggesting that SCA can identify groups of residues that act as evolutionary, structural, and functional sectors in a protein. It is important to note that these aspects can exist independently of one another. As an example, the existence of a physical inhomogeneity overlapping the statistical sector positions would support the idea that SCA can identify structural sectors, but would provide no guarantee that these also have an associated phenotype. For this reason, independent experimental verification is needed to support each of these claims.

We focus here on the experimental evidence supporting the hypothesis that SCA sectors act as functional sectors of proteins [8, 9, 18, 20, 24, 25]. We note that with the exception of Halabi et al. [18], this data refers to proteins in which a single SCA sector was identified, and we show that in this case, within statistical uncertainties, a method based on sequence conservation can identify functional residues as well as SCA. We also give a simple mathematical argument describing why this might happen.

Given that conservation information is explicitly used in calculating the SCA matrix, it is not surprising that SCA sectors are related to conservation. However, what we show here is that conservation dominates the SCA calculations in the single-sector case; thus, in order to establish whether the functional significance of SCA sectors is more than what is expected from single-site statistics alone, experiments need to focus on the examples where SCA identifies several sectors. The analysis of serine proteases described above provides such a study [18], but it is essential to have more data for different protein families to assess the robustness and generality of these observations.

2 Results

We analyze existing experimental datasets to compare the functional significance of SCA residues to that of conserved residues. Some of these datasets (PDZ, DHFR, and the voltage-sensing domains of potassium channels) have already been analyzed using SCA; one of them (*lacI*) has not. We show that in all these cases, conservation identifies functional positions just as effectively as SCA. This holds true for a wide range of choices of thresholds used to define conserved and functional

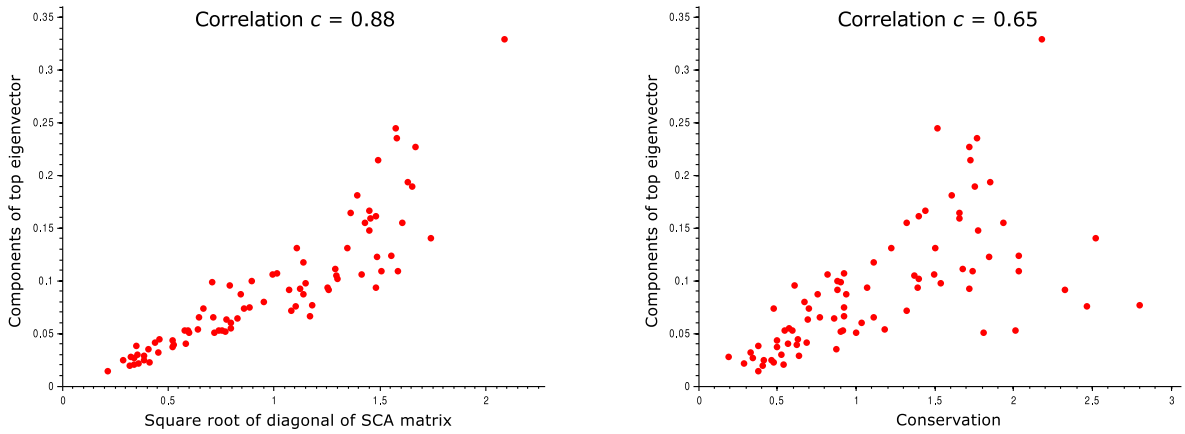
Correlation c = 0.88
Components of top eigenvector
Square root of diagonal of SCA matrix
Correlation c = 0.65
Components of top eigenvector
Conservation

Figure 1. Relation between the components of the top eigenvector of the SCA matrix and the square root of its diagonal elements *(left)* or conservation *(right)*. This was obtained for the PDZ alignment, but the results are similar for other alignments.

residues, respectively.

There have been several versions of the SCA approach that have been used in the literature. Indeed, each of the three datasets mentioned above for which SCA has been applied was analyzed using a different variation of the method. To avoid ambiguities, here we use a uniform method for all the alignments (see Methods for details). While Halabi et al. explicitly ignored the top eigenvector based on an analogy to finance [18], here we focus only on the top eigenvector. The reason for this is that, besides Halabi et al., all other published studies related to SCA have included this mode in their analysis [20, 24, 25].[1] The alignments used were generated using the HHblits software [30] with a consistent set of options (see Methods for details).

The definition of conservation we use is based on the relative entropy (Kullback-Leibler divergence) [18],

$$D_i = \sum_a f_i(a) \log \frac{f_i(a)}{q(a)} \,, \tag{1}$$

where $f_i(a)$ is the frequency at which amino acid a occurs in column i of the multiple sequence alignment and $q(a)$ is the background frequency for amino acid a. We use the same background frequencies as employed in SCA, which were calculated by Lockless and Ranganathan by averaging over a large protein database [13]. Other common definitions for conservation, such as the frequency of the most prevalent amino acid at a given position, tend to be highly correlated with D_i described above.

It is important to note that the qualitative results are unchanged regardless which version of SCA, which definition for conservation, or which alignments are used (see appendices for details). This work was motivated by the empirical observation that for many alignments the components of the top eigenvector correlate strongly with the diagonal elements of the SCA matrix (see Figure 1, left). The values of the diagonal elements can be calculated in terms of single-site statistics, raising the question whether correlations are needed to find the positions comprising the top sector. In fact, the components of the top eigenvector are also well-correlated with the conservation D_i defined

[1] In the case of DHFR, Reynolds et al. define the single sector using not only the top eigenvector, but the top five [24]. As we discuss below, the results from that paper are, however, not significantly changed if the sector is defined based only on the top eigenvector.

a) sector positions b) conserved residues

Figure 2. Contingency tables testing whether a PSD95^{pdz3} residue belonging to a sector or being highly-conserved is associated with significant functional effect upon mutagenesis. The tables are identical although only 57% of the residues are shared between the sector and the conserved positions.

above (see Figure 1, right). Note that these observations are not particularly surprising, given that the SCA matrix is weighted by quantities related to conservation. However, they also raise the question whether the observed functional significance of SCA sectors [24, 25, 31] could be due to conservation instead of correlations.

2.1 PDZ domains

The ability of SCA to identify residues that are important for protein function was recently tested in a high-throughput experiment involving a PDZ domain [25]. Each amino acid of the PSD95^{pdz3} domain was mutated to all 19 alternatives and the binding affinity of the resulting mutants to the PSD95^{pdz3} cognate ligand was measured. The measurement involved a bacterial two-hybrid system in which the PDZ domain was fused to the DNA-binding domain of the λ-cI repressor, while the ligand was fused to the α subunit of the *E. coli* RNA polymerase. This was used to control expression of GFP, which allowed the binding affinity between PSD95^{pdz3} and its ligand to be estimated using fluorescence-activated cell sorting (FACS). In order to quantify the sensitivity to mutations at a given site, the mutational effects on binding affinity were averaged over all 20 possible amino acids at that site. While mutations at most sites were found to have a negligible effect on ligand binding, 20 sites were identified where mutations had a significant deleterious effect [25].

The sector identified by SCA according to the methodology outlined above was found to indeed contain residues that are more likely to have functional significance than randomly chosen positions in the protein: 14 of the 21 sector residues are functionally significant, or 67%, compared to 25% for the entire protein (see Figure 2a). This is statistically-significant according to a Fisher exact test (one-tailed $p = 1 \times 10^{-6}$), and this result is robust to changing the threshold used to define the sector. This mirrors the results from McLaughlin Jr. et al. [25], obtained there with a different alignment constructed using a structural alignment algorithm [13].

There is another way of assessing the functional relevance of the sector positions that avoids making a sharp distinction between functional and non-functional residues [25]. The functional effects of mutations at all the positions in the domain were used to define a background distribution showing how likely an effect of a given magnitude was. If the sector is able to identify functionally-relevant positions, then the distribution of functional effects restricted to the sector positions should differ from this background distribution. Figure 3 (top) shows the comparison for the PDZ experiment. A two-sample Mann-Whitney U test [32] finds that indeed sector positions have a statistically-significant distribution of functional effects compared to all residues.

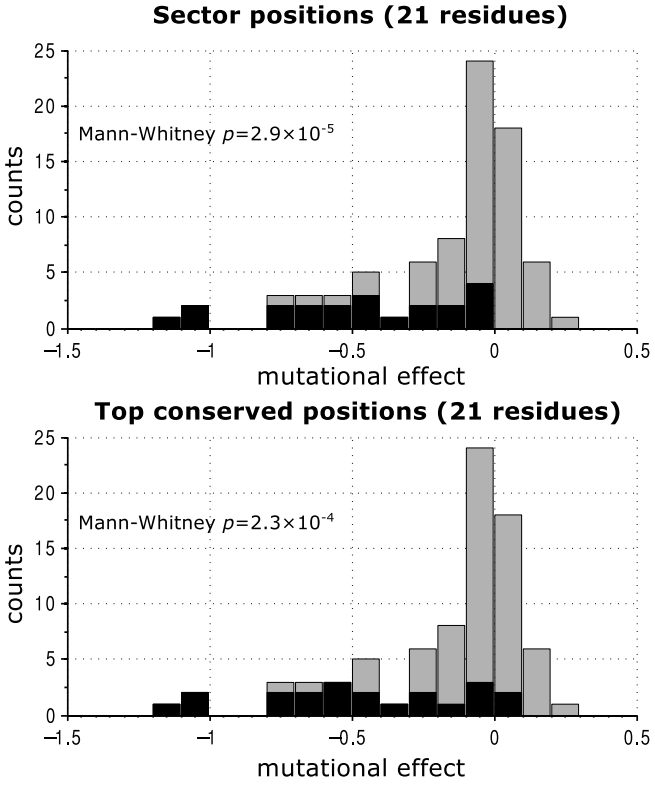

Figure 3. Histograms showing the effect of mutations on binding affinity of $PSD95^{pdz3}$ with cognate ligand, for all mutations (gray), and for mutations to selected positions (black). Each of the histograms in black contains 21 positions, with the largest SCA scores (top), or the largest conservation levels (bottom). A Mann-Whitney U test cannot find a statistically-significant difference between the distribution of mutational effects for sector positions and the one for conserved positions ($p = 0.9$). The mutational effect $\langle \Delta E_i^x \rangle_x$ is a dimensionless quantity calculated as in McLaughlin Jr. et al. [25].

We now test whether we could have obtained similar results by considering only sequence conservation. Indeed, although the 21 most conserved residues are different from the 21 residues identified by SCA (only about 60% are shared), the fraction of these residues that is functionally significant is the same (see Figure 2b). The histogram of functional effects is also essentially the same between SCA sector residues and conserved residues (see Figure 3, bottom), and in fact a Mann-Whitney U test confirms that the difference is not statistically significant.

McLaughlin Jr. et al. performed a similar analysis and obtained similar histograms as our Figure 3 (see Figure 3a in their paper [25]). The reason why our results seem so different is that, due to an error, the top histogram in Figure 3a in McLaughlin Jr. et al. is missing the data for the five sector residues that do not have a significant mutational effect. These five sector residues are mentioned and taken into account in other parts of the paper by McLaughlin Jr. et al [25], for example in Table S6a in the supplementary information, but they do not appear in the histogram. Had they been included, the histograms for conserved residues and that for SCA sector residues would look almost identical, in agreement with our results.

We stress again that these results do not imply that correlations in protein alignments are

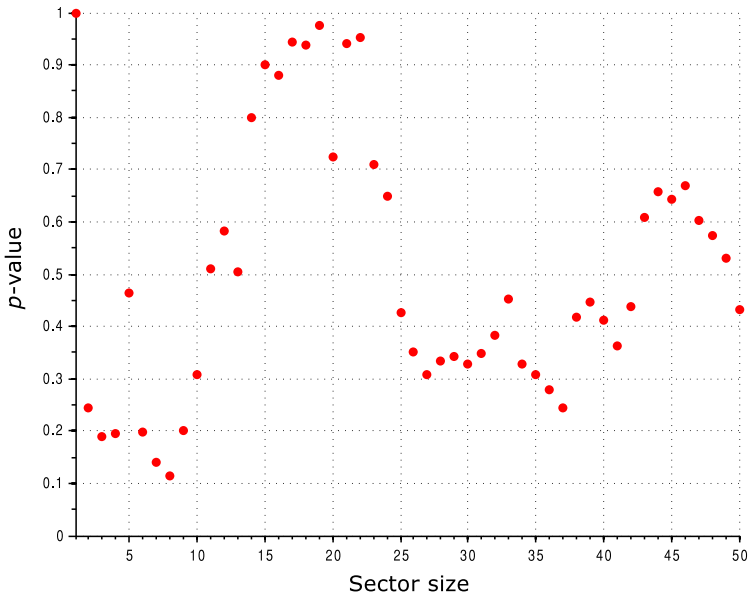

Figure 4. Comparison of the ability of the SCA sector and conservation to predict the functional effect of mutation of $PSD95^{pdz3}$ **residues for various sector sizes.** The vertical axis shows the p-value for a two-sample, two-tailed Mann-Whitney U test comparing the distribution of mutational effects for sector residues *vs.* conserved residues.

not informative. Indeed, as mentioned in the introduction, experimental data on the creation of artificial WW domains showed that ignoring correlations leads to non-functional proteins, while proteins designed based on conservation-weighted correlations can often be functional [8]. Moreover, correlation information was used to provide quite accurate predictions for contact maps and three-dimensional structures of a variety of proteins [10–12]. This is not possible using single-site statistics alone. The question we are asking, however, is whether the particular way in which alignment correlations are used in SCA is more useful for predicting functional information than conservation. The answer seems to be negative for the case of PDZ.

All the observations reported above are qualitatively the same when using different alignments, including the alignment employed by McLaughlin Jr. et al. [25] and a Pfam alignment. The observations are also robust to varying the threshold used for defining the sector: in Figure 4 we show a statistical comparison between the SCA sector and conserved residues calculated for various sizes of the sector.

Note that there are some potential caveats for the statistical tests we used. One assumption of both the Mann-Whitney U test and the χ^2 test employed above is that the samples analyzed are independent. In our case, the samples are the mutational effects at different residues in a protein domain, which are unlikely to be independent. Designing a statistical test that overcomes this difficulty would require a detailed model of evolutionary dynamics that accurately describes the relation between the binding affinity of $PSD95^{pdz3}$ to its cognate ligand, and the evolutionary information contained in a multiple sequence alignment. To our knowledge, there is unfortunately no unambiguous way of constructing such a model. Despite these issues, the analysis presented here suggests that, for the top sector, SCA is not significantly better than conservation at predicting functionally-important sites.

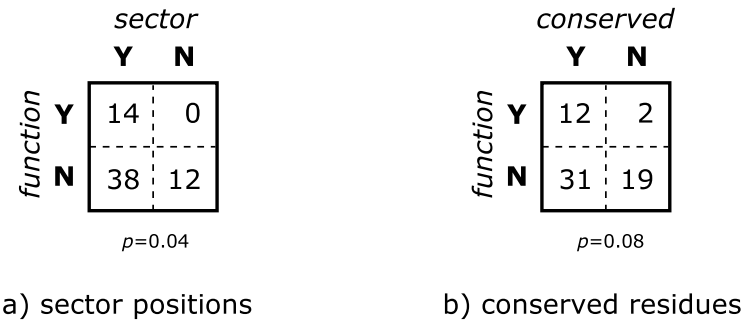

Figure 5. Contingency tables testing whether sector residues or conserved residues are more likely to "touch" functionally-significant LOV2 insertion points for the DHFR protein analyzed in Reynolds et al. [24]. A χ^2 test cannot reject the hypothesis that the two contingency tables are drawn from the same distribution ($p = 0.2$).

2.2 Dihydrofolate reductase

The case of dihydrofolate reductase (DHFR) [24] exhibits some interesting differences from PDZ. The experimental assay in this case involved perturbing the DHFR protein by attaching a light-sensitive domain (LOV2) between the atoms of the peptide bond immediately preceding each surface residue. The experiment used a folate auxotroph mutant of *E. coli* whose growth was rescued by a plasmid containing DHFR and thymidylate synthetase genes. The growth rate of the bacteria, which was measured with a high-throughput sequencing method, was shown to be approximately proportional to the catalytic efficiency of DHFR. The functional effect of each insertion of the LOV2 domain was measured by the difference in growth rates between lit and dark conditions. Out of the 61 measured surface sites, 14 were found to have a significant functional effect [24].

The effects of the insertion of the LOV2 domain are not localized on a single residue of the protein, which makes the analysis of the functional significance of the SCA sector positions more complicated in the case of DHFR. We follow here the method employed in the original study by Reynolds et al., which is to define a range around the insertion point within which a residue could conceivably feel the influence of the inserted domain [24]. More specifically, 4 Å spheres were centered on each of the four atoms forming the peptide bond broken by the insertion of LOV2, and any residues having at least one atom centered within any of these spheres was counted as "touching" the light-sensitive residue. The exact size of the cutoff is not important: we repeated the analysis with the cutoff set to 3 Å and 5 Å and obtained the same qualitative results.

Using the methodology described above, the SCA sector identified from the top eigenvector of the SCA matrix is found to "touch" all 14 of the functionally-significant LOV2 insertion sites. A set of conserved residues of the same size as the SCA sector "touches" 12 of the functionally-significant sites, and the difference is not statistically significant (see Figure 5).

The results we obtained for DHFR are somewhat less robust than those obtained for the other proteins. For the HHblits DHFR alignment, the qualitative result was the same for all sector sizes we tested (see Figure 6), but when using the Pfam alignment, very small SCA sectors (less than 10 residues) "touched" many more functionally-significant sites than sets of conserved residues of the same size. It is hard to verify whether this is a chance occurrence or a real phenomenon, and it is unclear whether the notion of a sector still makes sense when it comprises such a small part of the protein. One complication arises from the fact that highly conserved residues tend to cluster closer to the core of the protein (see Figure 7), and thus are less likely to "touch" its surface.

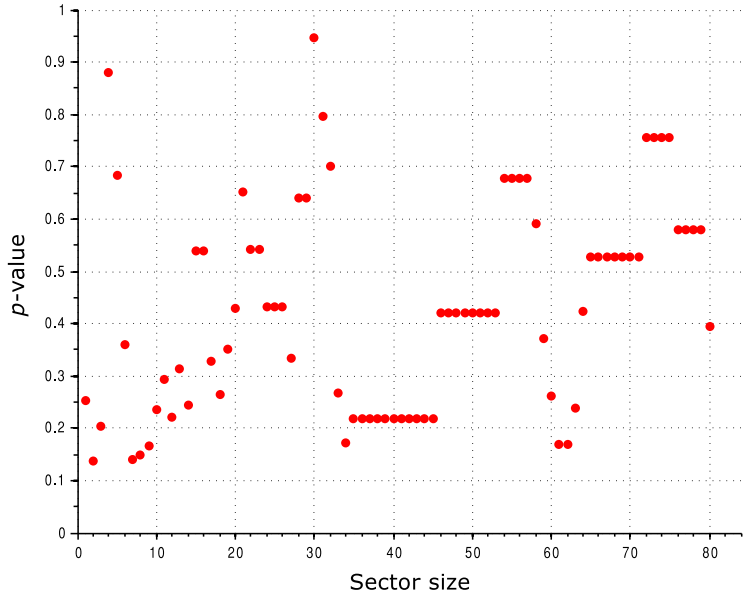

Figure 6. Comparison of the ability of the SCA sector and conserved residues to "touch" the functionally-significant sites of DHFR identified by Reynolds et al. [24]. The vertical axis shows the p-value for a two-tailed χ^2 test comparing the contingency tables obtained for the sector and for conservation (cf. Figure 5).

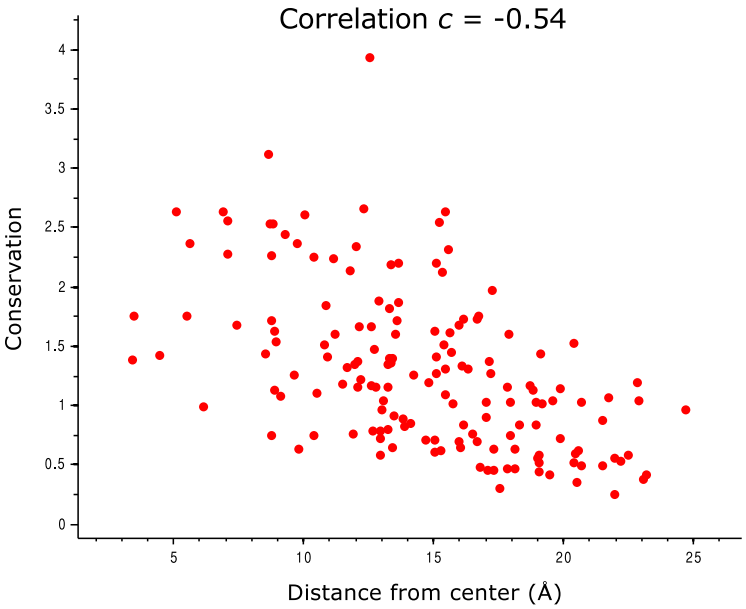

Figure 7. Dependence of conservation level on distance from the center of mass of DHFR protein.

a) sector positions b) conserved residues

Figure 8. Contingency tables testing whether belonging to a sector or being highly conserved is associated with significant functional effect upon mutagenesis for an alignment of voltage-sensing domains of potassium channels. Experimental data from Li-Smerin et al. [33]. The two contingency tables are identical although less than 80% of the residues are common between the SCA sector and the conserved positions.

2.3 Voltage-sensing domains of K^+ channels

Another dataset on which some work related to SCA has already been performed [31] was collected by Li-Smerin et al. [33]. In their experiments, 127 residues of the *drk1* K^+ channel were analyzed. For each of the mutants, voltage-activation curves were measured and fit to a two-state model, from which the difference in free energy between open and closed states ΔG_0 was estimated.

Following Lee et al. [31], we identified a set of functional sites using the condition $\left|\Delta G_0^{\mathrm{mut}} - \Delta G_0^{\mathrm{wt}}\right| \geq 1\,\mathrm{kcal/mol}$ and we compared this set to the SCA sector and to the most conserved residues. As with the other datasets, SCA and conservation turned out to be just as good at identifying functional positions in the voltage-sensing domains of potassium channels (see Figure 8).

2.4 *E. coli lac* repressor

A similar dataset to the PDZ dataset described above is available for the *lac* repressor protein in E. coli [34]. The authors used amber mutations and nonsense suppressor tRNAs to perform a comprehensive mutagenesis study of *lacI*. In this study, each one of 328 positions was mutated to 12 or 13 alternative amino acids, and the ability of each mutant protein to repress expression of the *lac* genes was tested. We summarized this data by recording, for each position, how many of the tested mutations had a significant effect on the phenotype of the *lac* repressor. We further identified "functionally-significant" sites by considering all the positions for which at least 8 substitutions resulted in loss of function. This threshold can be varied in the whole range from 1 to 10 without significantly altering the results.

As before, we observed a significant association between SCA sector positions and functional positions in the *lac* repressor; see Figures 9a and 10 (top). However, again, the set of most conserved positions was equally good at predicting functional sites—see Figures 9b and 10 (bottom). The results were not significantly affected by changing the size of the sector (see Figure 11).

2.5 Top eigenmode of the SCA matrix

In the previous sections, we showed that a significant fraction of the sector positions obtained from the top eigenvector of the SCA matrix can be predicted from single-site statistics. This can be attributed to a strong correlation between the components of the top eigenvector and the square root of the diagonal elements of the SCA matrix (see Figure 1, left). In Halabi et al., the top eigenvector

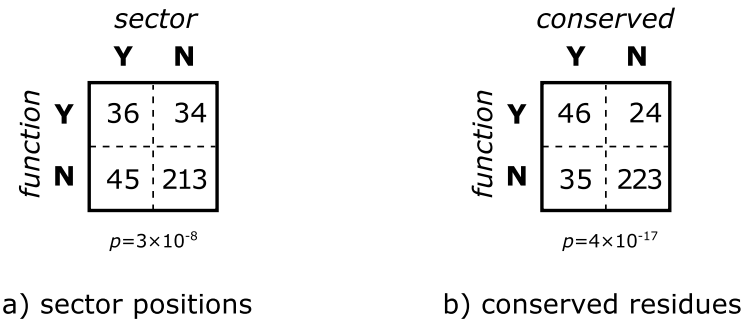

Figure 9. Contingency tables testing whether a *lac* repressor residue belonging to a sector or being highly-conserved is associated with significant functional effect upon mutagenesis. There is about 67% overlap between the two sets of residues. A two-tailed χ^2 test cannot reject the hypothesis that the two tables are drawn from the same distribution ($p = 0.2$).

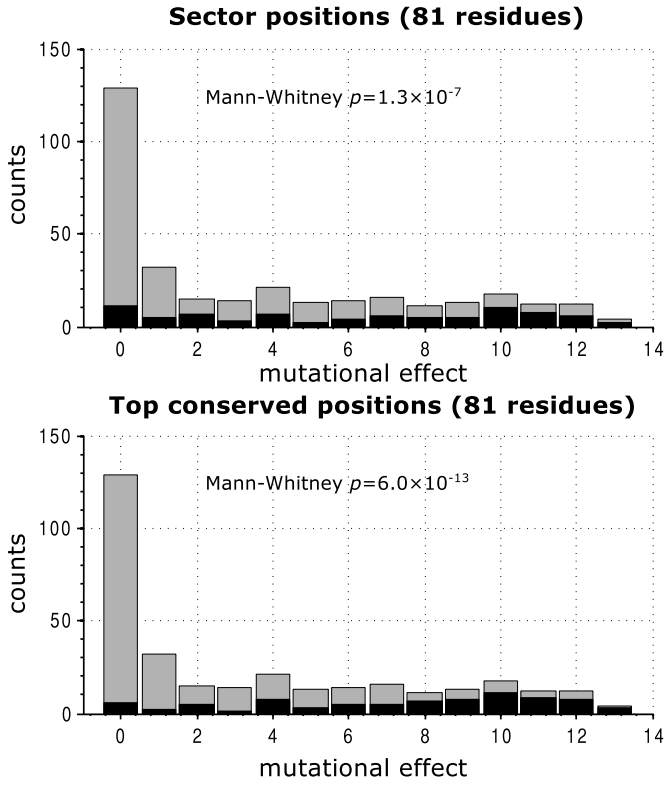

Figure 10. Histograms showing the effect of mutations on repression ability of *lacI* for all mutations (gray), and for mutations to selected positions (black). Each of the histograms in black contains 82 positions, with the largest SCA scores (top), or the largest conservation levels (bottom). While a Mann-Whitney U test finds the difference between the distribution of mutational effects for sector positions and the one for conserved positions bordering on statistical significance ($p \approx 0.08$), note that it is conservation that better matches the functional data.

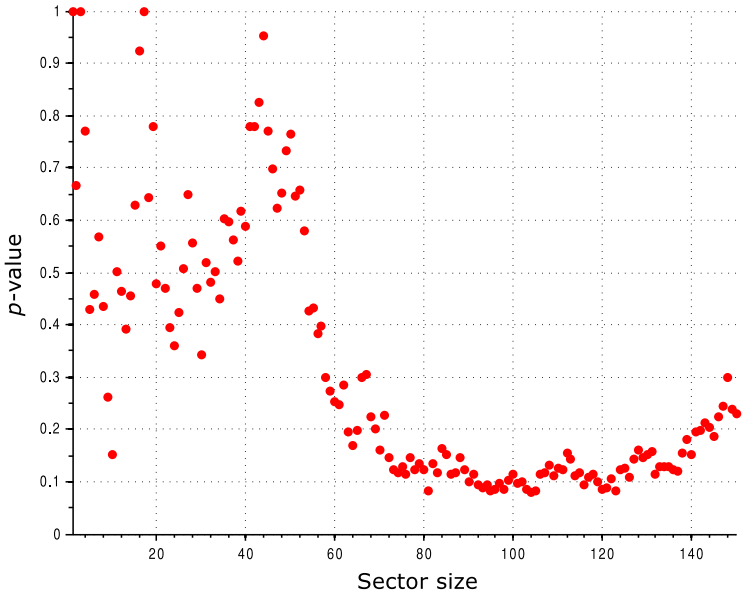

Figure 11. Comparison of the ability of the SCA sector and conservation to predict the functional effect of mutation of *lac* repressor residues for various sector sizes. The vertical axis shows the p-value for a two-sample, two-tailed Mann-Whitney U test comparing the distribution of mutational effects for sector residues *vs.* conserved residues. At large sector sizes, where the p value hovers around 0.1, it is the conserved residues that better match the functional data, rather than the SCA sector residues.

of the SCA matrix was ignored by analogy to finance, where this mode is a consequence of global trends in the market that affect all the stocks in the same way [18]. For proteins, the analogy is suggested to be with parts of sequences that are conserved due to phylogenetic relationships between the sequences in the alignment. Here we show that there is a different mechanism that can generate a spurious top eigenmode of the SCA matrix even when there are no phylogenetic connections between the sequences in the alignment. The main ingredient in this mechanism is a positive bias for the components of the SCA matrix.

Suppose that the underlying evolutionary process has no correlations between positions. Due to sampling noise, empirical correlations will typically be non-zero, and will fluctuate in a certain range. We denote the size of these fluctuations by x. The off-diagonal elements of the covariance matrix will have mean zero and variances of order $C_{ij}^2 \sim C_{ii}C_{jj}x^2$. In this case, the reason for the positive bias for the components of the SCA matrix is the fact that typically SCA takes the absolute value of the covariances (or some norm that produces only non-negative values; see appendix B) [18,24,25]. This implies that the off-diagonal entries of this matrix will have expectation values of order $x\sqrt{C_{ii}C_{jj}}$.[2]

Even when the absolute value is not used, the correlation between the components of the top eigenmode of the SCA matrix and the diagonal elements of this matrix may also occur; this happens for example for the alignment in Smock et al. [20]. Simulations involving random alignments show that this phenomenon occurs whenever there are weak, uniform correlations between all the positions in an alignment. This can be the result of phylogenetic bias, but could have a different origin. This situation could be distinguished from the one above by looking at how the magnitude x of the off-diagonal correlations scales with alignment size; it should scale roughly like the inverse of the number of sequences if it is due to sampling noise, and be approximately constant otherwise.[3]

To try to explain these empirical observations, let us consider a simplified version of the SCA matrix:

$$M = \begin{pmatrix} \Delta_1 & d_1 d_2 x & \cdots & d_1 d_n x \\ d_2 d_1 x & \Delta_2 & \cdots & d_2 d_n x \\ \cdots & \cdots & \ddots & \cdots \\ d_n d_1 x & d_n d_2 x & \cdots & \Delta_n \end{pmatrix} . \tag{2}$$

Writing out the eigenvalue equation and performing some simple algebraic manipulations reveals that the eigenvector components v_i corresponding to eigenvalue λ are related to the diagonal elements Δ_i by

$$\frac{\sqrt{\Delta_i}}{v_i} \propto \lambda - \frac{\Delta_i}{1+x} . \tag{3}$$

When the top eigenvalue is much larger than the other ones, which is usually the case when applying SCA to protein alignments, the following approximation holds:

$$\lambda_{\text{top}} \approx \frac{x}{1+x} \sum_i \Delta_i . \tag{4}$$

Empirically, this is observed to roughly match the results of SCA on real protein alignments. Given that $\lambda_{\text{top}} \gg \Delta_i$, we can also write

$$v_{i,\text{top}} \approx \frac{\alpha}{\lambda_{\text{top}}} \times \sqrt{\Delta_i} , \tag{5}$$

where α is a normalization constant. This is the observed linear relation between the top eigenvector and the square root of the diagonal elements of the SCA matrix (Figure 1, left). Note that the

[2] Note that the positional weights can be absorbed into the diagonal elements C_{ii}, so we do not write them out explicitly.

[3] We thank D. Hekstra for this observation.

SCA matrix for an alignment does not really have the highly symmetric form (2); instead it shows fluctuations in the off-diagonal components. Because of this, we cannot expect to see all the eigenvectors obey eq. (3). Indeed, for SCA matrices obtained from protein alignments, eq. (3) seems to hold only for the top eigenvector. A treatment of this problem in the framework of random matrix theory might help to clear up the expectations one should have for the top eigenvector of the SCA matrix, but such an analysis goes beyond the scope of this paper.

The simple argument described above suggests that, under certain conditions that seem to hold in the cases where SCA has been applied, the top eigenvector of the SCA matrix is indeed related to conservation, and is largely independent of correlations between positions. This does not mean that there is no information contained in this top mode, but does imply that most of this information can be obtained by looking at single-site statistics alone.

Note again that in our derivation the origin of the off-diagonal entries is not specified. While we showed that they can be a simple artifact of sampling noise, they could also be partly due to a non-trivial phylogenetic structure of the alignment, as previously suggested [18].

3 Discussion

3.1 Proteins with multiple SCA sectors

It is perhaps not surprising that conservation is a good indicator of the functionally-important residues in a protein; indeed, this fact is one of the original motivations for using positional weights in SCA that grow with conservation levels [19]. However, as a consequence, for proteins with a single SCA sector, it is difficult to distinguish between the functional significance of sector residues and that of conserved residues. The natural solution to this problem is to focus on proteins with multiple sectors, such as the serine protease family analyzed by Halabi et al. [18].

In the serine protease case, three SCA sectors were identified by placing thresholds on certain linear combinations of eigenvectors of the SCA matrix.[4] These sectors (called 'blue', 'red', and 'green') were found to have independent effects on various phenotypes of the protein: the blue sector affected denaturation temperature, the red one affected binding affinity, and the green sector contained the residues responsible for catalytic activity.

There are two attractive features of the serine protease data. One is that several different quantities were measured for each mutant, thus allowing for a test of the idea that the protein is split into groups each of which affects different phenotypes. Another important feature is that some double mutants were also measured, showing that mutations in different sectors act approximately independently from each other. Collecting more extensive data of this type for serine proteases and for other proteins should give more weight to the idea that SCA sectors act as functional sectors in proteins. To reduce the amount of work involved, we point out that from our observations, it seems that instead of a complete scan of all 19 alternative amino acids at each position, an alanine scan, involving only mutations to alanine, might be sufficient. Using only alanine replacements, even a complete double-mutant study of $\mathrm{PSD95}^{\mathrm{pdz3}}$ would require about 3000 mutants, only a factor of two more than were already studied [25]. For proteins exhibiting multiple SCA sectors, this number could be lowered by focusing only on those double mutants that combine mutations in different sectors, thus testing the independence property.

Finding several relevant quantities to measure for each of the mutants might not be an easy task. An ideal system for this would be related to gene expression or signal transduction, allowing

[4] The top eigenvector was ignored based on an analogy to finance, and thus the issues outlined in the previous section do not apply here.

measurements to be made in realistic conditions. Furthermore, it would be convenient to have a low-dimensional quantitative description of the protein's phenotype, so that one could check whether the sectors predicted by SCA correlate with the mutations that affect the parameters in this description.

One difficulty in the application of SCA is that the identification of sectors is non-trivial. Halabi et al. used visual inspection to identify linear combinations of eigenvectors to represent the sectors [18]. Independent component analysis (ICA) has also been invoked to find the linear combinations [19, 20, 22], but a mathematically rigorous motivation for the application of this procedure is missing. An approach that avoids these difficulties is to check whether a linear regression can approximate the measured quantities for the different mutants with linear combinations of the eigenvectors of the SCA matrix. This seems to work for the case of serine protease (see Figure S1), though the small number of data points prevents a statistically rigorous analysis. A similar approach does not work for the PDZ data from McLaughlin Jr. et al., in which binding to both the cognate (CRIPT) ligand and to a mutated $T_{-2}F$ ligand was measured [25] (see Figure S2). It also does not work for the potassium channels dataset, in which both the activation voltage V_{50} and the equivalent charge z were measured for each mutant [33] (see Figure S3). This is consistent with the idea that these proteins exhibit a single sector.

Conservation alone cannot in general be used to find several distinct groups of residues that have distinct functions. For this reason, finding evidence for functionally significant and independent SCA sectors would automatically favor SCA over a simple conservation analysis. However, it is important to point out that SCA, with the particular set of weights as defined by Halabi et al. [18], is only one possible procedure for analyzing correlations in sequence alignments. Once more data is available for proteins containing multiple sectors, it will be important to compare different sets of positional weights, or different models altogether, to identify the best approach for analyzing MSAs [23].

3.2 Conclusions

We analyzed the available evidence regarding the hypothesis that the residues comprising the sectors identified by statistical coupling analysis are functionally significant. We looked at a number of studies, some directly related to SCA [18, 24, 25], and some unrelated [33, 34], and we showed that while the sector positions identified by SCA do tend to be functionally relevant, in the case of single-sector proteins, conserved positions provide a statistically equivalent match to the experimental data. This observation was traced to a property of the SCA matrix that makes the components of its top eigenvector correlate strongly with its diagonal entries. We presented a mathematical model that might explain this correlation. This model suggests that, as a generic property of statistical coupling analysis, the top eigenvector of the SCA matrix does not contain information beyond that provided by single-site statistics.

The observation that conservation is an important determinant of the SCA sectors is of course not very surprising, since one of the principles of SCA is to upweight the correlation information for conserved residues compared to poorly-conserved ones. However, this does pose a problem for the interpretation of the large-scale experiments that have been performed in relation to SCA [24, 25], given that these provide most of the available evidence for the functional significance of SCA sectors. Our analysis show that this functional significance might be due to conservation alone. Since function is not the only reason for which protein residues may be conserved [35], it is not surprising that the overlap with functional residues is not perfect.

Once again, it is important to note that our findings do not imply that correlations within MSAs are uninformative; the contrary seems to be supported by experimental data [8, 10–12]. However,

in order to test whether the particular way in which these correlations are used within the SCA framework is useful for making functional predictions about proteins, it will be necessary to go beyond single-sector proteins and measure several different phenotypes. Such data exists [18], but is too limited at this point to be conclusive. A thorough investigation of the idea that SCA sectors act as functional sectors requires more of this type of data, for a wider class of proteins.

Whether small groups of residues inside proteins act as independent "knobs" controlling the various phenotypes is a question that can be asked independently of any statistical analysis of alignments. Such functional sectors could be found by mutagenesis work, as described above. Alternatively, one could look for structural sectors using NMR or X-ray data to search for correlated motions. This has the advantage of not requiring the modification of proteins through mutations. Finally, evolutionary sectors could be searched for by using artificial evolution experiments. If the existence of these functional, structural, or evolutionary sectors is verified with sufficient precision, one could then more easily approach the question of whether a statistical method is capable of inferring their composition from an MSA, and in this case, which method is the most efficient and accurate.

4 Methods

4.1 Sequence alignments

Statistical coupling analysis requires an alignment of protein sequence homologs as input data. This may contain both orthologs and paralogs, and at least moderate sequence diversity within the alignment is necessary, because an alignment of identical sequences will not contain any information about amino acid covariance. The alignments we used were generated using HHblits, with an E-value of $E = 10^{-10}$. States with 40% or more gaps were considered insert states, and were later removed from the calculations. The Uniprot IDs of the seed sequences used with HHblits are as follows: DLG4_RAT (PDZ), DYR_ECOLI (DHFR), KCNB1_RAT (K^+ channels), and LACI_ECOLI (*lacI*). To check the robustness of the results, we also ran our analysis on Pfam alignments when available, and on the alignments from McLaughlin Jr. et al. [25], Reynolds et al. [24], and from Lee et al. [31] for the PDZ, DHFR, and potassium channels datasets, respectively.

4.2 Statistical coupling analysis

The statistical coupling analysis was performed in accordance with the projection method [19, 25], which is the default in the newest version of the SCA framework from the Ranganathan lab. The code we used for the analysis can be accessed at **https://bitbucket.org/ttesileanu/multicov**.

Consider a multiple sequence alignment represented as an $N \times n$ matrix A in which a_{ki} is the amino acid at position i in the k^{th} sequence. We first construct a numeric matrix $\tilde{X}$ by

$$\tilde{x}_{ki} = \begin{cases} \dfrac{\phi_i(a_{ki}) f_i(a_{ki})}{\sqrt{\sum\limits_{b \neq \text{gap}} \phi_i^2(b) f_i^2(b)}} & \text{if } a_{ki} \neq \text{gap}\,, \\ 0 & \text{if } a_{ki} = \text{gap}\,, \end{cases} \tag{6}$$

where $\phi_i(a)$ is a positional weight, and $f_i(a)$ the frequency with which amino acid a occurs in column i of the alignment. The positional weights are given by

$$\phi_i(a) = \log\left[\frac{f_i(a)}{1 - f_i(a)} \frac{1 - q(a)}{q(a)}\right], \tag{7}$$

where $q(a)$ is the background frequency with which amino acid a occurs in a large protein database. The SCA matrix is, up to an absolute value, the covariance matrix associated with $\tilde{X}$,

$$\tilde{C}_{ij} = \left| \frac{1}{N} \sum_k \tilde{x}_{ki} \tilde{x}_{kj} - \frac{1}{N^2} \sum_{k,l} \tilde{x}_{ki} \tilde{x}_{lj} \right| . \tag{8}$$

Finally, the sector was identified by finding the positions where the components of the top eigenvector of $\tilde{C}_{ij}$ went above a given threshold. The threshold was chosen so that the sector comprised about 25% of the number n of residues contained in each alignment sequence.

More details about this method and descriptions of the other variants of SCA found in the literature can be found in appendix B.

4.3 Sequence conservation

The conservation level of a position in the alignment is calculated using the relative entropy (Kullback-Leibler divergence), as described in eq. (1). A different definition, as the frequency of the most prevalent amino acid at a position, is highly correlated with D_i and gives similar results.

Note that the calculation of the relative entropy as defined in eq. (1) requires that $\sum_a f_i(a) = 1$ and $\sum_a q(a) = 1$. For the first of these relations to hold, we need the sum over a to include the gap, but this requires a value for the background frequency of gaps $q(\text{gap})$. This is not straightforward to estimate or even to define. There are several solutions possible: one is to assume that the background frequency for gaps is equal to the gap frequency in the alignment averaged over all positions. Another approach is to simply ignore the gaps by focusing only on the sequences that do not contain a gap at position i. We chose the former solution, as it is the default one in the SCA framework, but the results are very similar when using the latter choice.

5 Acknowledgments

We are grateful to Richard McLaughlin Jr., Rama Ranganathan, and Kim Reynolds for sharing their scripts and data with us, and for useful discussions. We would also like to thank Gérard Ben Arous, Doeke Hekstra, Michael Mitchell, Rama Ranganathan, and Olivier Rivoire for discussions and comments on early drafts of this manuscript. T. T. was supported by a Charles L. Brown Membership at the Institute for Advanced Study. L. J. C. was supported by an Engineering and Physical Sciences Research Council Fellowship (EP/H028064/2).

A Calculating sequence covariance

Let us assume we are given a multiple sequence alignment of homologous proteins as an $N \times n$ matrix A_{ki}. In order to calculate correlations in the alignment, we need a way of transforming it to numeric data. There are three main ways in which this was done in the literature [18, 20, 22, 24, 25], which we describe below; generally all these methods yield highly similar results. In the present work, we used the projection method (item 3 below), which is the default in the latest SCA release from the Ranganathan lab. Note that SCA requires positional weighting to be done on top of the covariance analysis. This is described in the next section.

1. **The binary approximation.**

   We start by describing the simplest approach, which constructs a "binary approximation" of the alignment [18, 20]. In the binary approximation, each amino acid is replaced by 1 if it is equal to the consensus amino acid at its position, and 0 otherwise,

   $$X_{ki} = \delta(A_{ki}, c_i)\,, \tag{S1}$$

   where the consensus amino acid c_i is the most frequent amino acid found in column i of the alignment.[5] Here A_{ki} is the amino acid found at position i in sequence k, and the Kronecker symbol $\delta(a, b)$ is 1 if and only if $a = b$.

   The covariance matrix is then defined in the standard way,

   $$C_{ij}^{\text{bin}} = f_{ij} - f_i f_j\,, \tag{S2}$$

   where

   $$f_{ij} = \frac{1}{N}\sum_k X_{ki} X_{kj}\,, \qquad f_i = \frac{1}{N}\sum_k X_{ki}\,, \tag{S3}$$

   with N being the number of sequences in the alignment. Note that f_i is simply the frequency of the consensus amino acid at position i, which is often used as a measure of the conservation level at that position. Also, f_{ij} is the frequency at which the consensus amino acids occur simultaneously in the two columns i and j.

2. **The reduction method.**

   The most generic statistical analysis that can be performed with categorical data is using *contingency tables*. In this context, these are tables of the frequencies at which various combinations of amino acids occur simultaneously in a sequence—we can, for example, define the frequency $f_i(a)$ with which amino acid a is found at position i, and the frequency $f_{ij}(a, b)$ with which amino acids a and b co-occur at positions i and j, respectively. It is then convenient to define a "binary representation" of the alignment [22], $x_{ki}(a)$, where $x_{ki}(a)$ is equal to 1 if A_{ki} is a, and 0 otherwise; in short,

   $$x_{ki}(a) = \delta(A_{ki}, a)\,. \tag{S4}$$

   Note that here there are 21 columns for each column in the original alignment. It is important to keep in mind that—despite the potentially confusing nomenclature—the binary representation $x_{ki}(a)$ is very different from the binary approximation X_{ki}. The former is an exact

[5] This is typically restricted to non-gaps, though in practice this usually does not affect the results.

representation of the alignment data, while the latter is an approximation that keeps only part of the information in A_{ki}.

The single-site and pairwise frequencies, $f_i(a)$ and $f_{ij}(a,b)$, are thus averages involving the binary representation

$$\begin{aligned} f_i(a) &= \frac{1}{N}\sum_k x_{ki}(a)\,, \\ f_{ij}(a,b) &= \frac{1}{N}\sum_k x_{ki}(a)x_{kj}(b)\,. \end{aligned} \tag{S5}$$

Now we can define the covariance

$$C_{ij}(a,b) = f_{ij}(a,b) - f_i(a)f_j(b)\,. \tag{S6}$$

For each pair of sites, we obtain a covariance value between every pair of amino acids, a and b. There are 21 choices of amino acid (including the gap), so for each pair of sites (i,j), C_{ij} is a 21×21 matrix. Note, however, that frequencies are normalized—$\sum_a f_i(a) = 1$ and $\sum_b f_{ij}(a,b) = f_i(a)$—which means that not all numbers in the covariance matrix are independent. In fact, if we are given the values of $C_{ij}(a,b)$ for all amino acids except one (commonly, the gap), it is straightforward to infer the missing values. For this reason, unless otherwise stated, we will always assume a and b to exclude the gap, and will treat the matrix C_{ij} as having size 20×20 instead of 21×21.

In the binary approximation, for each pair of sites i and j, we were able to calculate one number C_{ij}^{bin} showing the amount of covariance between the sites. In contrast, the full covariance matrix $C_{ij}(a,b)$ contains an entire matrix of numbers for each pair of sites. It is sometimes useful to collapse this matrix to a single number that measures the overall covariance between two sites, as was the case in the binary approximation. This is usually done using a heuristic approach, for example by using the Frobenius norm [19]

$$C_{ij}^{\text{red}} = \left[\sum_{a,b} C_{ij}^2(a,b)\right]^{1/2}, \tag{S7}$$

or the spectral norm [24]

$$C_{ij}^{\text{red}} = \text{largest singular value of } C_{ij}(a,b)\,. \tag{S8}$$

The binary approximation described above can actually also be seen as a reduction method, in which

$$C_{ij}^{\text{red}} = C_{ij}(c_i,c_j)\,, \tag{S9}$$

with c_i being the consensus amino acid in column i, as before.

Another approach for carrying out the reduction starts from a ratio of frequencies,

$$D_{ij}(a,b) = \log\frac{f_{ij}(a,b)}{f_i(a)f_j(b)} = \log\left[1+\frac{C_{ij}(a,b)}{f_i(a)f_j(b)}\right]. \tag{S10}$$

The mutual information is then a natural measure of the independence of two random variables that can be constructed from D [36, 37],

$$\mathrm{MI}_{ij} = \sum_{\substack{a,b \\ \text{incl. gaps}}} f_{ij}(a,b) D_{ij}(a,b) = \sum_{\substack{a,b \\ \text{incl. gaps}}} f_{ij}(a,b) \log \frac{f_{ij}(a,b)}{f_i(a) f_j(b)} \,. \tag{S11}$$

Note that here we include gaps in the summation because the definition of the mutual information given in eqs. (S10) and (S11) assumes that $f_i(a)$ and $f_{ij}(a,b)$ are proper frequencies; in particular, they should sum up to 1, and this is only true if the values for the gaps are included in the sum.

3. **The projection method.**

Instead of calculating the full covariance matrix and then reducing it, the projection method starts by projecting the binary representation onto an $N \times n$ numeric matrix Y_{ki},

$$Y_{ki} := \sum_a x_{ki}(a) v_i(a) \,, \tag{S12}$$

where v_i are unit vectors. We can then use the definitions (S2) and (S3) with Y instead of X to obtain the covariance matrix,

$$C_{ij}^{\mathrm{proj}} = \frac{1}{N} \sum_k Y_{ki} Y_{kj} - \frac{1}{N^2} \sum_{k,l} Y_{ki} Y_{lj} \,. \tag{S13}$$

The projection vectors that we use in this paper are given by [25]:

$$v_i(a) = \frac{f_i(a)}{\sqrt{\sum_b f_i(b)^2}} \,. \tag{S14}$$

Note that, as we did above for the covariance matrix, here too we do not include the gap, *i.e.,* we assume that $v_i(\mathrm{gap}) = 0$, and we have $b \neq \mathrm{gap}$ in the summation.

It is interesting to note that a different choice for the v_i can be used to recover the binary approximation,

$$v_i^{\mathrm{binary}}(a) = \delta\big(a, \arg\max_{b \neq \mathrm{gap}} f_i(b)\big) \equiv \delta(a, c_i) \,. \tag{S15}$$

For both the binary approximation and the projection method, the approximation is most accurate for highly conserved sites.

B Statistical coupling analysis

There are two main features that distinguish SCA from a standard covariance analysis: the introduction of positional weights and the positivity of the matrix elements. Starting again with the binary approximation, for simplicity, the covariance matrix C defined in (S2) is transformed to [18]

$$\tilde{C}_{ij} = |\phi_i \phi_j C_{ij}| \,, \tag{S16}$$

where the weights ϕ_i are specific functions of the single-site statistics, chosen based on the idea that entries of the covariance matrix corresponding to poorly conserved sites are less likely to be informative than those corresponding to highly conserved sites, and thus should be given less weight.

The absolute value that appears in the formula was justified by the requirement of finding blocks of coevolving residues regardless of the sign of the correlations [18]. It also avoids a certain instability to small perturbations that appears in cases in which the consensus amino acid has a frequency that is very similar to that of the next most common amino acid. In these cases, a small perturbation in the alignment can flip the order of the top amino acids, thus flipping the sign of some correlations in the binary approximation.[6] Taking the absolute value may, however, introduce artifacts, as described in the main text.

The expression commonly used for the positional weights is [18, 20, 25]

$$\phi_i = \log\left[\frac{f_i}{1-f_i}\frac{1-q(c_i)}{q(c_i)}\right], \tag{S17}$$

where f_i is the conservation at site i, *i.e.* the frequency in the alignment of the consensus amino acid, while $q(c_i)$ are background frequencies for the consensus amino acids c_i. Note that we are assuming that the background frequencies depend only on the identity of the consensus amino acid, and not on the position i in the protein. These background expectations can be estimated by averaging over a large set of proteins. The functional form (S17) for the positional weights was chosen to match the original 1999 formulation of SCA [13, 18] and to fulfill the role of down-weighting poorly conserved sites, but is otherwise arbitrary.

Instead of performing the positional weighting at the level of the covariance matrix, it could have been performed on the binary alignment itself. The covariance matrix of the transformed binary alignment

$$\tilde{X}_{ki} = \phi_i X_{ki} \tag{S18}$$

directly yields $\tilde{C}_{ij}$, after taking the absolute value of each element. This can be generalized to apply to the binary representation matrix; we define

$$\tilde{x}_{ki}(a) = \phi_i(a) x_{ki}(a)\,, \tag{S19}$$

in which we replace (S17) by [22]

$$\phi_i(a) = \log\left[\frac{f_i(a)}{1-f_i(a)}\frac{1-q(a)}{q(a)}\right]. \tag{S20}$$

With this positionally-weighted binary representation we can use either the reduction method [22] or the projection method [25] described in the previous section. When using the projection method, the absolute value of each element is taken when calculating the SCA matrix, as is done with the binary approximation. An interesting empirical observation is that the mutual information defined in eq. (S11) is well-approximated by a weighted SCA matrix using the reduction method, in which the positional weights are chosen equal to the logarithm of the frequencies.

For the results in the main text, after the SCA matrix was calculated, the sector positions were identified by placing a threshold on the components of its top eigenvector. Note that, due to the positivity of the elements of the SCA matrix, the Perron-Frobenius theorem guarantees that all these components are positive.

[6] We thank O. Rivoire for this observation.

C Alignments

As described in the Methods section of the paper, the software package HHblits [30] was used to generate the alignments used in the the main text. We also analyzed a number of alignments from other sources to make sure the results weren't sensitive to this choice. For the cases of PDZ [25], DHFR [24], and the potassium channels [31], we ran the analysis on the alignments that were used in the papers that first applied the SCA method to those proteins. For PDZ and DHFR we also ran the analysis on Pfam alignments (PF00595 and PF00186, respectively). This was not done for the potassium channels and *lacI* because no suitable alignments were available in Pfam. See Table S1 for a summary of the results.

To improve the quality of the alignments, it has been suggested [22] to filter them by removing repeated sequences, sequences with too many gaps, and positions with too many gaps. From our tests, however, these procedures do not make a big difference to the results of this paper. For this reason, the only filtering we perform is to remove the insert states from HHblits alignments, which in our case amounts to removing columns containing 40% or more gaps.

In addition, for the PDZ alignment from McLaughlin Jr. et al. [25], we removed the columns containing more than 20% gaps, because this is how that alignment was processed in the scripts provided by the authors of that article [25]. There is a minor glitch in the procedure of mapping alignment columns to PDB coordinates in McLaughlin Jr. et al. that leads to the misidentification of one of the columns (corresponding to PDB position 334). For consistency with the older work, we worked with this minor error in the alignment, but we checked that the results are not significantly affected by it.

D Details about the DHFR analysis

The way in which the DHFR alignment was analyzed by Reynolds et al. [24] has a number of peculiarities compared to the other datasets we presented, which we describe below. We note, however, that their results are not significantly different from those obtained with our simplified protocol.

The SCA method used by Reynolds et al. was the spectral-norm reduction method described above (see section A, item 2) using a thresholded form of the positional weights [24],

$$\phi_i(a) = \begin{cases} \log\left[\dfrac{f_i(a)}{1-f_i(a)}\dfrac{1-q_a}{q_a}\right] & \text{for } f_i(a) > q_a, \\ 0 & \text{else.} \end{cases} \tag{S21}$$

Furthermore, the SCA matrix was "cleaned" by subtracting the average SCA matrix calculated for 100 randomized alignments. Each of the randomized alignments was obtained by independently permuting the elements of the alignment columns, which has the effect of destroying correlations without affecting the single-site amino acid frequencies.[7]

The sector was defined by the residues for which the component of at least one of the top five eigenvectors goes above a given threshold [24]. To select the threshold, first a Student's t-distribution was fit to the components of each of the eigenvectors, and then the value for which the t-distribution PDF drops below a certain threshold was used as a cutoff. The PDF threshold is given by p_i for the

[7] Since the number of random samples is finite, the results depend slightly on the state of the random number generator. The results from the original paper by Reynolds et al. can be obtained by using the default random number generator in Matlab with the default seed.

Table S1. Results of the analysis performed in this paper when run on different alignments.

Alignment	Contingency table for sector[a]			Contingency table for conservation			Comparison (χ^2 p value)
		S	**NS**		**C**	**NC**	
PDZ [25]	**F**	15	5	**F**	15	5	$p_{\chi^2} = 1.00$
	NF	8	53	**NF**	8	53	
		S	**NS**		**C**	**NC**	
PDZ (Pfam)	**F**	13	7	**F**	9	11	$p_{\chi^2} = 0.45$
	NF	7	52	**NF**	11	48	
		S	**NS**		**C**	**NC**	
DHFR[b] [24]	**F**	14	0	**F**	12	2	$p_{\chi^2} = 0.30$
	NF	33	17	**NF**	27	23	
		S	**NS**		**C**	**NC**	
DHFR[c] (Pfam)	**F**	14	0	**F**	12	2	$p_{\chi^2} = 0.29$
	NF	35	15	**NF**	29	21	
		S	**NS**		**C**	**NC**	
potassium channels[d] [31]	**F**	18	19	**F**	18	19	$p_{\chi^2} = 0.96$
	NF	17	67	**NF**	20	64	

[a] **C**, **S**, and **F** stand for conserved, sector, and functional, respectively.

[b] For DHFR we are counting the residues that are "touched" by either sector or conserved residues. Although the number of conserved residues we are considering is equal to the number of residues in the sector, the number of surface residues that are "touched" is different.

[c] See footnote *b*.

[d] The number of residues in the SCA sector is the same as the number of conserved residues we considered. The sum of the entries in the **S** column does not match that in the **C** column because some of the sector residues are located at sites where we have no experimental data.

Table S2. The results from the paper when restricting to alanine mutations.

Alignment	Contingency table for sector			Contingency table for conservation			Comparison (χ^2 p value)
PDZ		**S**	**NS**		**C**	**NC**	$p_{\chi^2} = 0.98$
	F	10	10	**F**	9	11	
	NF	11	50	**NF**	12	49	
lacI		**S**	**NS**		**C**	**NC**	$p_{\chi^2} = 0.82$
	F	16	26	**F**	20	22	
	NF	65	221	**NF**	61	225	

i^{th} eigenvector, where p_i is[8]

$$p_i = \frac{0.005}{2\,\mathrm{IQR}(v_i)n^{-1/3}}\,, \tag{S22}$$

where $\mathrm{IQR}(v)$ is the interquartile range of v (the range over which the middle half of the components of v spread), and v_i is the i^{th} eigenvector. Despite the complicated selection procedure, a very similar sector can be obtained by using a constant PDF threshold for each of the top five eigenvectors, or even just by using the largest components of only the top eigenvector.

E Alanine scans

Instead of using all the experimental data for PDZ and for *lacI*, we can restrict our attention to alanine mutations, to get an idea for the amount of information contained in an alanine scan. As mentioned in the paper, the qualitative results do not change much, though, as expected, the quality of the match between the predictions from SCA or conservation and the experimental data is reduced (see Table S2).

F Diagonal of SCA matrix instead of conservation

In the paper we point out that the top eigenvector of the SCA matrix correlates primarily with the diagonal elements $\tilde{C}_{ii}$ of this matrix (or rather, with their square root), while the correlation with conservation is weaker. This is because, although both conservation and the diagonal elements of the SCA matrix can be calculated from the single-site frequencies $f_i(a)$, the relation between them is non-trivial and non-monotonic. We therefore wondered how the functional significance of the sector residues compared to that of residues that have high values of $\tilde{C}_{ii}$. The results can be found in Table S3.

[8]The dependence of the threshold on i is an artifact of the fact that Reynolds et al. applied a constant threshold to histogram values instead of PDF values, and the bin size for the histograms was determined using the Freedman-Diaconis rule [38], and thus varied between eigenvectors. This can be seen from the Matlab scripts provided by the authors.

Table S3. The results of the paper when replacing conservation by the diagonal of the SCA matrix.

Alignment	Contingency table using $\tilde{C}_{ii}$			Comparison to sector (χ^2 test p value)
		C	**NC**	
PDZ	**F**	12	8	$p_{\chi^2} = 0.87$
	NF	9	52	
		C	**NC**	
DHFR	**F**	13	1	$p_{\chi^2} = 0.61$
	NF	34	16	
		C	**NC**	
potassium channels	**F**	16	21	$p_{\chi^2} = 0.95$
	NF	16	74	
		C	**NC**	
lacI	**F**	43	27	$p_{\chi^2} = 0.55$
	NF	38	220	

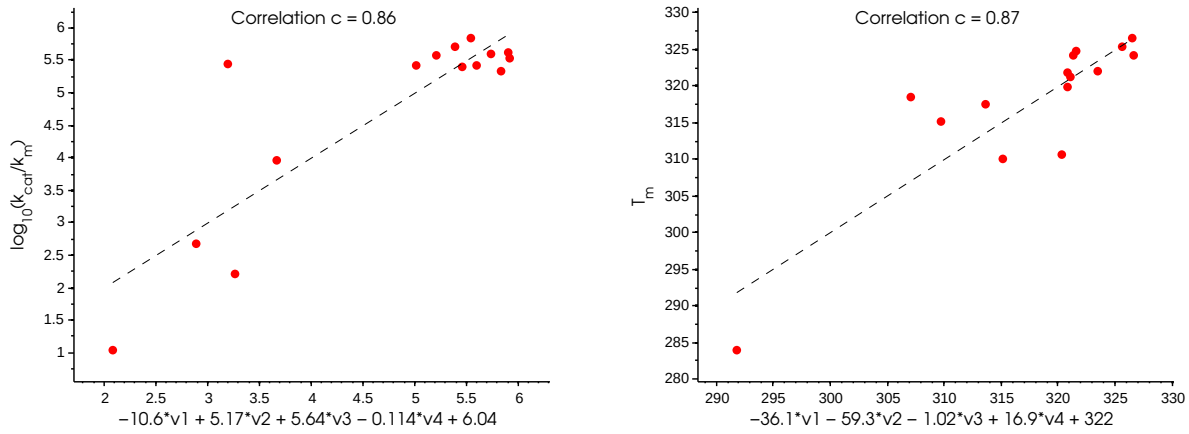

Figure S1. We attempt to fit binding affinity (left) or denaturation temperature (right) for the single mutants of rat trypsin described in Halabi et al. [18] against the components of the top four eigenvectors of the SCA matrix corresponding to the mutated residues. The best linear regressions are shown on the x-axis. The dashed line has slope 1 and intercept 0.

G Multiple sectors

There are two key questions related to multiple SCA sectors: one is how to determine how many eigenvectors to analyze, and the second one is which linear combinations of eigenvectors to use for finding sectors. We briefly discussed the second question in the main text, and pointed out that while there are several approaches that have been used in the literature to find the appropriate linear combinations, these approaches have little or no theoretical motivation and have been tested only in a very limited fashion.

In the main text, we pointed out that we can instead use linear regression to find the linear combinations that best approximate the measured quantities.[9] Here we show the results of such an analysis for the case of serine protease (Figure S1), PDZ (Figure S2), and the potassium channels (Figure S3). While for serine protease the regression works well for both measured quantities, for PDZ we can only fit the mutational effect on binding to the CRIPT ligand, and for potassium channels the fit is not very good to either the activation potential V_{50} or the equivalent charge z.

The number of eigenvectors to consider for sector determination is itself a difficult problem. It was addressed by Halabi et al. using an approach inspired from the analysis of financial markets [18]. Essentially, the eigenvalue spectrum of the SCA matrices obtained for randomized alignments was obtained, and a threshold was established at the upper edge of this distribution. Any eigenvalues of the SCA matrix for the real alignment that fell below this edge were considered noise and were ignored. This approach was motivated by the Marčenko-Pastur distribution, that describes the eigenvalue spectrum for a covariance matrix $C = X^T X$ associated to a random i.i.d. data matrix X. However, the assumption of identically distributed elements does not hold for alignment data, because different columns in the alignment have different amino acid distributions, even when there are no correlations between columns. For this reason, the spectrum of the SCA matrix does not resemble the Marčenko-Pastur distribution, and intuitions based on this distribution need not hold in the case of SCA. In particular, eigenvalues that are below the edge obtained from randomized

[9] This of course assumes that the relation is linear, which is far from obvious, but can be thought of as a first-order approximation.

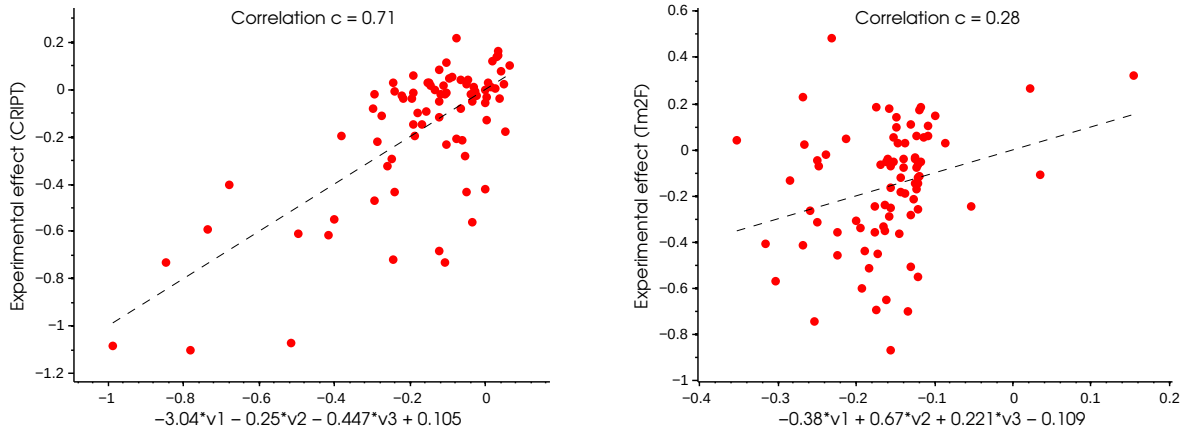
Correlation c = 0.71
Experimental effect (CRIPT)
−3.04*v1 − 0.25*v2 − 0.447*v3 + 0.105
Correlation c = 0.28
Experimental effect (Tm2P)
−0.38*v1 + 0.67*v2 + 0.221*v3 − 0.109

Figure S2. We attempt to fit the measured mutational effect for binding to the CRIPT ligand (left) or the $T_{-2}F$ ligand (right) as measured for the single mutants of $PSD95^{pdz3}$ described in McLaughlin Jr. et al. [25] against the components of the top three eigenvectors of the SCA matrix corresponding to the mutated residues. The best linear regressions are shown on the x-axis. The dashed line has slope 1 and intercept 0.

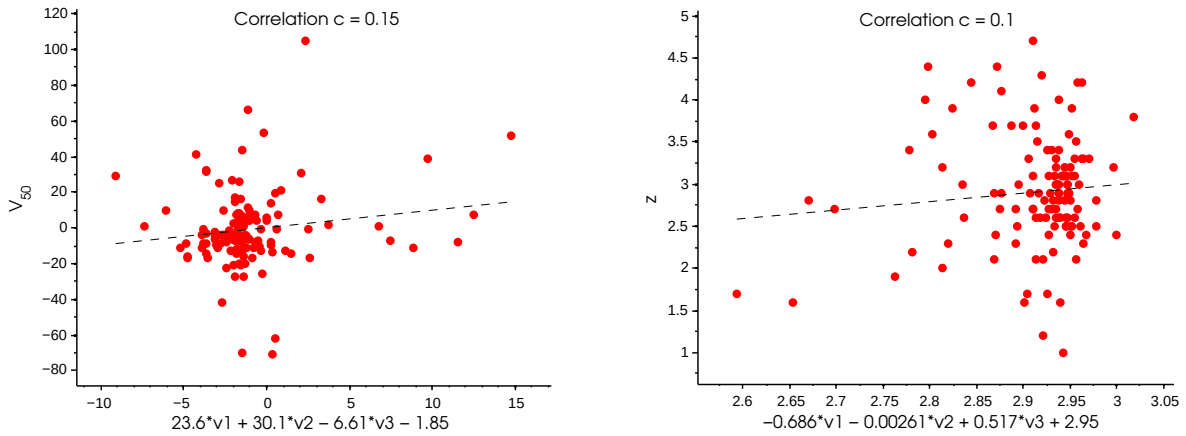

Figure S3. We attempt to fit the activation voltage V_{50} (left) and the equivalent charge z (right) measured for single mutants of the *drk1* voltage-gated K^+ channel in Li-Smerin et al. [33] against the components of the top three eigenvectors of the SCA matrix corresponding to the mutated residues. The best linear regressions are shown on the x-axis. The dashed line has slope 1 and intercept 0.

samples could contain perfectly valid information about the protein.

Another difficulty with SCA is that the absolute value of all the elements in the covariance matrix is taken, which makes the top eigenvector become an outlier. As described in the main text, although this top eigenmode is much over the "noise" edge described above, the information it contains is essentially independent of correlations, and is thus indistinguishable from noise.[10]

H Top eigenmode of SCA matrix—some details

Here we fill in some of the details for the model presented in the paper that can explain the correlation between the components of the top eigenvector of the SCA matrix and its diagonal entries. As in the paper, suppose we have a covariance matrix with off-diagonal entries that are biased towards positive values. The simple model we wrote for this is

$$M = \begin{pmatrix} d_1^2(1+x) & d_1 d_2 x & \cdots & d_1 d_n x \\ d_2 d_1 x & d_2^2(1+x) & \cdots & d_2 d_n x \\ \cdots & \cdots & \ddots & \cdots \\ d_n d_1 x & d_n d_2 x & \cdots & d_n^2(1+x) \end{pmatrix} \equiv \begin{pmatrix} \Delta_1 & d_1 d_2 x & \cdots & d_1 d_n x \\ d_2 d_1 x & \Delta_2 & \cdots & d_2 d_n x \\ \cdots & \cdots & \ddots & \cdots \\ d_n d_1 x & d_n d_2 x & \cdots & \Delta_n \end{pmatrix} . \tag{S23}$$

For simplicity, let us assume that there are no degeneracies between the d_i, *i.e.,* that $d_i \neq d_j$ for $i \neq j$, and that M is not singular, *i.e.,* $d_i \neq 0$ for all i. Let $v = (v_1, \ldots, v_n)$ be an eigenvector of this matrix with eigenvalue λ. Then we have

$$d_i^2 v_i + d_i x \sum_j d_j v_j \stackrel{!}{=} \lambda v_i \,, \tag{S24}$$

which yields[11]

$$v_i = \Big(\sum_j d_j v_j \Big) \frac{d_i x}{\lambda - d_i^2} \,. \tag{S25}$$

This implies that the components of the eigenvectors are related to the diagonal elements $\Delta_i = d_i^2(1+x)$ by

$$\frac{\sqrt{\Delta_i}}{v_i} \propto \lambda - \frac{\Delta_i}{1+x} \,. \tag{S26}$$

If we multiply eq. (S25) by d_i and sum over i, we can divide through by $\sum_j d_j v_j$, and get

$$1 = x \sum_i \frac{d_i^2}{\lambda - d_i^2} \,, \tag{S27}$$

which can be used to estimate λ. In particular, this equation allows us to show that between each consecutive pair of values d_{i_1} and d_{i_2}, there is exactly one eigenvalue.

[10] This is because the randomized alignments used with SCA, which are used to calculate the noise distribution, are obtained by shuffling elements in the alignment columns. This keeps the single-site frequencies intact while destroying correlations, implying that the single-site frequencies are a feature of the noise model.

[11] We may worry about division by zero. Note that, from the eigenvalue equation (S24), $\lambda = d_i^2$ for some i if and only if $\sum_j d_j v_j = 0$ (since we assumed all $d_i \neq 0$). However, feeding this back into eq. (S24), we see that this is only possible if all the v_j for which $d_j^2 \neq \lambda$ are zero. Since we assumed that none of the d_j vanish, $\sum_j d_j v_j = 0$ can only hold if at least two components of v are non-vanishing. This, however, would imply that there is a degeneracy, which we explicitly disallowed. We thus conclude that $\lambda \neq d_i^2$ for any i.

By the Perron-Frobenius theorem, the top eigenvector can be chosen to have all components positive, and thus it should have λ larger than all d_i^2. Assuming $\lambda \gg d_i^2$, which empirically seems to be the case for SCA matrices, we get an estimate for the top eigenvalue

$$\lambda_{\text{top}} \approx x \sum_i d_i^2 = \frac{x}{1+x} \sum_i \Delta_i \equiv \frac{x}{1+x} \operatorname{Tr} M \,. \tag{S28}$$

It should be checked that this is consistent with the condition that λ is much larger than all d_i^2; this seems to be true for empirical SCA matrices. The Perron-Frobenius theorem also guarantees that all other eigenvectors of M will have both positive and negative components, and therefore, according to eq. (S25), the corresponding eigenvalues will have to be smaller than the largest d_i^2. This implies that for the SCA matrices, the top eigenvector will be an outlier, *i.e.,* the SCA matrices are approximately rank-1, which can indeed be observed for real alignments.

Using $\lambda_{\text{top}} \gg d_i^2$ in eq. (S25), we get

$$v_{i,\text{top}} \approx \Big(\sum_j d_j v_{j,\text{top}}\Big) \frac{x}{\lambda_{\text{top}}} \times \sqrt{\Delta_i} \,, \tag{S29}$$

which is the observed linear relation between the top eigenvector and the square root of the diagonal elements of the SCA matrix. This argument shows that the top eigenvector is strongly correlated with single-site statistics and thus largely independent of correlations between positions. It is important to emphasize that this does not mean that there is no information contained in this mode, but only that most of this information can be obtained without any analysis of correlations.

As mentioned in the paper, we emphasize again that in this derivation the origin of the off-diagonal entries is not specified. They could be an artifact of sampling noise, they could come from actual non-specific correlations between positions, or they could be due to a non-trivial phylogenetic structure of the alignment, as suggested by Halabi et al. [18].

I Data availability

This study is based on several alignments that are either directly accessible online or upon request from the authors of earlier studies. Some of the alignments can also be created using publicly-available software. This is described in detail below:

1. HHblits alignments for PDZ, DHFR, *lacI*, and the K^+ channels were generated as described in sections 4.1 and C. HHblits can be obtained from `ftp://toolkit.genzentrum.lmu.de/pub/HH-suite/`.

2. Pfam alignments were downloaded from `http://pfam.xfam.org/`.

3. The PDZ and DHFR alignments used in McLaughlin Jr. et al. 2012 [25] and Reynolds et al. 2011 [24], respectively, were obtained upon request from the Ranganathan lab, `http://systems.swmed.edu/rr_lab/index.html`.

4. The potassium channels alignment from Lee et al. 2009 [31] is available as supporting information to that article, `http://www.plosbiology.org/article/info:doi/10.1371/journal.pbio.1000047#s5`.

The mutagenesis data that we used is available from the following sources:

1. The PDZ mutation data is available directly from the Ranganathan lab website, `http://systems.swmed.edu/rr_lab/papers/McLaughlin_Ranganathan/McLaughlin_Ranganathan_data.mat.zip`.
2. The DHFR light-sensitivity data was obtained upon request from the Ranganathan lab.
3. The K^+ channels data we used is available in the main text of Li-Smerin et al. 2000 [33].
4. The *lacI* data from Markiewicz et al. 1994 [34] is available for download at `http://sift.bii.a-star.edu.sg/www/test_sets/LacITable.txt`.

The raw alignment and mutagenesis data were further processed using custom-made Matlab and C++ code, publicly available for download from `https://bitbucket.org/ttesileanu/multicov`.

J Differences from previous version

There are several differences between this version and the previous one uploaded to the arXiv. The most important difference is that we decided to use a consistent SCA method and a consistent methodology for generating the alignments for all the datasets we consider.[12] Previously we were using the SCA method and alignments from the earlier studies that had compared SCA to mutational data. We also fixed a bug that caused a small error in the conservation values.

[12] This was suggested to us by a reviewer.

References

1. Do CB, Katoh K (2008) Protein multiple sequence alignment. Methods in Molecular Biology 484: 379–413.

2. Notredame C (2002) Recent progress in multiple sequence alignment: a survey. Pharmacogenomics 3: 131–144.

3. Schneider TD, Stormo GD, Gold L, Ehrenfeucht A (1986) Information content of binding sites on nucleotide sequences. Journal of Molecular Biology 188: 415–431.

4. Zvelebil MJ, Barton GJ, Taylor WR, Sternberg MJE (1987) Prediction of protein secondary structure and active sites using the alignment of homologous sequences. Journal of Molecular Biology 195: 957–961.

5. Hollstein M, Sidransky D, Vogelstein B, Harris CC (1991) p53 Mutations in Human Cancers. Science 253: 49–53.

6. Cargill M, Altshuler D, Ireland J, Sklar P, Ardlie K, et al. (1999) Characterization of single-nucleotide polymorphisms in coding regions of human genes. Nature Genetics 22: 231–238.

7. Ng PC, Henikoff S (2003) SIFT: predicting amino acid changes that affect protein function. Nucleic Acids Research 31: 3812–3814.

8. Russ WP, Lowery DM, Mishra P, Yaffe MB, Ranganathan R (2005) Natural-like function in artificial WW domains. Nature 437: 579–583.

9. Socolich M, Lockless SW, Russ WP, Lee H, Gardner KH, et al. (2005) Evolutionary information for specifying a protein fold. Nature 437: 512–518.

10. Marks DS, Colwell LJ, Sheridan R, Hopf TA, Pagnani A, et al. (2011) Protein 3D Structure Computed from Evolutionary Sequence Variation. PLoS ONE 6: e28766.

11. Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, et al. (2011) Direct-coupling analysis of residue co-evolution captures native contacts across many protein families. PNAS 108: E1293–E1301.

12. Weigt M, White RA, Szurmant H, Hoch JA, Hwa T (2009) Identification of direct residue contacts in protein-protein interaction by message passing. PNAS 106: 67–72.

13. Lockless SW, Ranganathan R (1999) Evolutionarily Conserved Pathways of Energetic Connectivity in Protein Families. Science 286: 295–299.

14. Fuentes EJ, Der CJ, Lee AL (2004) Ligand-dependent dynamics and intramolecular signaling in a PDZ domain. Journal of Molecular Biology 335: 1105–1115.

15. Peterson FC, Penkert RR, Volkman BF, Prehoda KE (2004) Cdc42 regulates the Par-6 PDZ domain through an allosteric CRIB-PDZ transition. Molecular Cell 13: 665–676.

16. Süel GM, Lockless SW, Wall MA, Ranganathan R (2003) Evolutionarily conserved networks of residues mediate allosteric communication in proteins. Nature 10: 59–69.

17. Hatley ME, Lockless SW, Gibson SK, Gilman AG, Ranganathan R (2003) Allosteric determinants in guanine nucleotide-binding proteins. PNAS 100: 14445–14450.

18. Halabi N, Rivoire O, Leibler S, Ranganathan R (2009) Protein sectors: evolutionary units of three-dimensional structure. Cell 138: 774–786.

19. Ranganathan R, Rivoire O (2012). Note 109: A summary of SCA calculations. Available online at `http://systems.swmed.edu/rr_lab/Note109_files/Note109_v3.pdf`.

20. Smock RG, Rivoire O, Russ WP, Swain JF, Leibler S, et al. (2010) An interdomain sector mediating allostery in Hsp70 molecular chaperones. Molecular Systems Biology 6: 414.

21. Bouchaud JP, Potters M (2009) Financial Applications of Random Matrix Theory: a short review. arXiv 0910.1205.

22. Rivoire O (2013) Elements of Coevolution in Biological Sequences. Physical Review Letters 110: 178102.

23. Colwell LJ, Brenner MP, Murray AW (2014) Conservation weighting functions enable covariance analyses to detect functionally important amino acids. PLoS ONE 9: e107723.

24. Reynolds KA, McLaughlin Jr RN, Ranganathan R (2011) Hot spots for allosteric regulation on protein surfaces. Cell 147: 1564–75.

25. McLaughlin Jr RN, Poelwijk FJ, Raman A, Gosal WS, Ranganathan R (2012) The spatial architecture of protein function and adaptation. Nature 491: 138–142.

26. Kern D, Zuiderweg ERP (2003) The role of dynamics in allosteric regulation. Current Opinion in Structural Biology 13: 748–757.

27. Matoba Y, Sugiyama M (2003) Atomic resolution structure of prokaryotic phospholipase A2: analysis of internal motion and implication for a catalytic mechanism. Proteins 51: 453–69.

28. Fraser JS, Clarkson MW, Degnan SC, Erion R, Kern D, et al. (2009) Hidden alternative structures of proline isomerase essential for catalysis. Nature 462: 669–673.

29. Dhulesia A, Gsponer J, Vendruscolo M (2008) Mapping of two networks of residues that exhibit structural and dynamical changes upon binding in a PDZ domain protein. Journal of the American Chemical Society 130: 8931–8939.

30. Remmert M, Biegert A, Hauser A, Söding J (2012) HHblits: lightning-fast iterative protein sequence searching by HMM-HMM alignment. Nature Methods 9: 173–175.

31. Lee SY, Banerjee A, MacKinnon R (2009) Two separate interfaces between the voltage sensor and pore are required for the function of voltage-dependent K(+) channels. PLoS Biology 7: 676–686.

32. Mann HB, Whitney DR (1947) On a test of whether one of two random variables is stochastically larger than the other. Annals of Mathematical Statistics 18: 50–60.

33. Li-Smerin Y, Hackos DH, Swartz KJ (2000) Alpha-helical structural elements within the voltage-sensing domains of a K(+) channel. The Journal of General Physiology 115: 33–49.

34. Markiewicz P, Kleina LG, Cruz C, Ehret S, Miller JH (1994) Genetic Studies of the lac Repressor – XIV. Analysis of 4000 Altered Escherichia coli lac Repressors Reveals Essential and Non-essential Residues, as well as "Spacers" which do not Require a Specific Sequence. Journal of Molecular Biology 240: 421–433.

35. Mirny LA, Shakhnovich EI (1999) Universally Conserved Positions in Protein Folds: Reading Evolutionary Signals about Stability, Folding Kinetics and Function. Journal of Molecular Biology 291: 177–196.

36. Fodor AA, Aldrich RW (2004) Influence of conservation on calculations of amino acid covariance in multiple sequence alignments. Proteins 56: 211–21.

37. Dunn SD, Wahl LM, Gloor GB (2008) Mutual information without the influence of phylogeny or entropy dramatically improves residue contact prediction. Bioinformatics 24: 333–40.

38. Freedman D, Diaconis P (1981) On the histogram as a density estimator: L2 theory. Zeitschrift fuer Wahrscheinlichkeitstheorie und Verwandte Gebiete 57: 453–476.